arXiv:1510.03661v1 [math.PR] 13 Oct 2015

Segregating Markov chains

Timo Hirscher * Anders Martinsson †
Chalmers University of Technology and University of Gothenburg

October 14, 2015

Abstract

Dealing with finite Markov chains in discrete time, the focus often lies on convergence behavior and one tries to make different copies of the chain meet as fast as possible and then stick together. There is, however, a very peculiar kind of discrete finite Markov chain, for which two copies started in different states can be coupled to meet almost surely in finite time, yet their distributions keep a total variation distance bounded away from 0, even in the limit as time goes off to infinity. We show that the supremum of total variation distance kept in this context is $\frac{1}{2}$.

1 Introduction

When the long time behavior of Markov chains is analyzed, one of the most common strategies is to couple several copies of the chain started in different states. In doing so, one standard approach is to define two copies of a Markov chain (started in different states) on a common probability space, correlated in some clever way such that they are likely to meet up within some moderate time, and glue them together as soon as they meet.

This idea is so predominant that little attention was directed away from such couplings; in the standard reference [4] it was even claimed erroneously that *any* coupling of two Markov chains with the same transition probabilities can be modified so that the two chains stay together at all times after their first simultaneous visit to a single state. A counterexample to this statement was in fact given in [5]: If a coupling of two copies of the same Markov chain is changed in such a way that the second copy mimicks the behavior of the first one once they meet, the altered individual process might no longer be a copy of the given chain.

For the sake of simplicity, we want to restrict our considerations to time-homogeneous Markov chains that evolve in discrete time and on countable state spaces – even if the argument used to derive Theorem 4.1 in fact applies to

*Research supported by grants from the Swedish Research Council and the Royal Swedish Academy of Sciences
†Research supported by grants from the Swedish Research Council

more general settings as well, see Remark 1. So let $\mathbf{X} = (X_n)_{n\in\mathbb{N}_0}$ denote a Markov chain on the countable state space S with transition probabilities $\{P(r,s) = \mathbb{P}(X_{n+1} = s \,|\, X_n = r);\ r,s \in S,\ n \in \mathbb{N}_0\}$. We will denote the distribution of X_n given $X_0 = x$, i.e. for a copy of the chain started in $x \in S$ at time n, by $P^n(x,.)$.

In what follows, we want to describe and analyze an exceptional kind of Markov chain, that was first mentioned and introduced by Häggström [3]: chains in which two copies started in different states can be coupled such that they almost surely meet, but their distributions do not come arbitrarily close to one another with respect to total variation distance. We will refer to this behavior as *segregation* of two states.

As a preparation, the second section deals with the concept of couplings and convergence of Markov chains. Much of this is very standard, but there is also the lesser known but crucial distinction between Markovian and faithful couplings. Section 3 introduces the idea of segregating Markov chains and tries to shed light on their peculiar nature.

In Section 4, more precisely in Theorem 4.1, we prove that the value $\frac{1}{2}$ is an upper bound on the total variation distance that the distributions of copies of such a Markov chain, started in two states that the chain segregates, can maintain.

In Section 5, a constructive and intuitively accessible example of a Markov chain is given, that segregates two states such that the total variation distance kept can be pushed arbitrarily close to $\frac{1}{\mathrm{e}}$. This improves on the example in [3] and serves as a warmup for the more technical and implicit construction in the last section, that further improves this value.

Finally, in Section 6 we introduce and employ the idea of *separation* to show that for any $\varepsilon > 0$ there exist Markov chains segregating states x and y such that copies started in these two states can be coupled to meet almost surely while their distributions $P^n(x,.)$ and $P^n(y,.)$ have a total variation distance of at least $\frac{1}{2} - \varepsilon$ for all $n \in \mathbb{N}$, see Theorem 6.2.

Together with the upper bound from Section 4, this answers the question of how large the supremum of total variation distances that can be kept in this context is and establishes our main result, Theorem 6.1, stating that $\frac{1}{2}$ is in fact the right answer.

2 Preliminaries: convergence and couplings

In order to quantify the difference between two probability measures (such as the distributions of two copies of a Markov chain at a fixed time) there are quite a few distance measures. The so-called total variation distance is among the most common ones.

Definition 1
Let μ and ν be two probability distributions on a countable set S. The *total*

variation distance between the two measures is then defined as

$$\|\mu - \nu\|_{\mathrm{TV}} := \sup_{A \subseteq S} |\mu(A) - \nu(A)|.$$

This notion of distance is used in most of the standard convergence theorems on finite Markov chains as well (e.g. see Thm. 4.9 in [4]):

Theorem 2.1
Suppose $(X_n)_{n\in\mathbb{N}_0}$ is an irreducible and aperiodic Markov chain on a finite state space S. Then there exists a unique limiting distribution π on S, called the stationary distribution, as well as constants $\alpha \in (0,1)$ and $C > 0$ such that

$$\|P^n(x,.) - \pi\|_{\mathrm{TV}} \leq C\,\alpha^n,$$

for all $x \in S,\ n > 0$.

If the distribution of a Markov chain at time n converges to the same distribution π as n tends to infinity, irrespectively of its starting distribution, the standard way to measure the speed of convergence is the following non-increasing function of n:

$$d(n) := \sup_{x\in S} \|P^n(x,.) - \pi\|_{\mathrm{TV}}.$$

Sometimes it is however more convenient to consider the related function

$$\overline{d}(n) := \sup_{x,y\in S} \|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}}.$$

On the one hand, this function is not only non-increasing but also submultiplicative, i.e. $\overline{d}(m+n) \leq \overline{d}(m)\cdot\overline{d}(n)$, on the other it has a close relationship to d in that for all $n \in \mathbb{N}_0 :\ d(n) \leq \overline{d}(n) \leq 2\,d(n)$, see Lemmas 4.11 and 4.12 in [4]. Note that the proofs there are for finite S but immediately transfer to our setting with S being countable.

On the basis of this notion of distance, the central concept of mixing time is defined, loosely speaking, as the time it takes until the effect of the starting distribution has begun to disappear substantially.

Definition 2
Given a Markov chain, for which the distribution of X_n converges to a fixed distribution π (irrespectively of the distribution of X_0), define its *mixing time* by

$$t_{\mathrm{mix}} := \min\{n \in \mathbb{N}_0;\ d(n) \leq \tfrac{1}{4}\}.$$

As already mentioned, the tool that often makes proofs about convergence of Markov chains both short and elegant is the coupling approach. Let us therefore properly define this standard construction and then highlight which additional properties a coupling can have.

Definition 3
We define a *coupling* of two copies of a Markov chain on S to be a process $((X_n, Y_n))_{n\in\mathbb{N}_0}$ on $S \times S$, with the property that both $(X_n)_{n\in\mathbb{N}_0}$ and $(Y_n)_{n\in\mathbb{N}_0}$ are Markov chains on S with the same transition probabilities (but possibly different starting distributions).

If the process $((X_n, Y_n))_{n\in\mathbb{N}_0}$ is itself a Markov chain (not necessarily time-homogeneous), it is called a *Markovian coupling*.

In order to get good estimates on mixing times it is often of importance to bring into line the long term behavior of the chain started in different states. In order to do so, one wants to make sure that the two coupled chains stay together once they meet, more precisely: if $X_m = Y_m$, then $X_n = Y_n$, for all $n \geq m$. Couplings with this property are sometimes called "*sticky*" couplings. As noted in the introduction, it is however not possible to modify every coupling in such a way that it becomes sticky by simply glueing together the two copies once they meet, see Prop. 3 in [5] for an example. The crucial property is the following:

Definition 4
A Markovian coupling $((X_n, Y_n))_{n\in\mathbb{N}_0}$ of two copies of a Markov chain is called *faithful* if for all $x_n, y_n, x_{n+1}, y_{n+1} \in S$, $n \in \mathbb{N}_0$:

$$\begin{aligned}\mathbb{P}\big(X_{n+1} = x_{n+1} \,|\, (X_n, Y_n) = (x_n, y_n)\big) &= \mathbb{P}(X_{n+1} = x_{n+1} \,|\, X_n = x_n) \\ &= P(x_n, x_{n+1})\end{aligned}$$

and

$$\begin{aligned}\mathbb{P}\big(Y_{n+1} = y_{n+1} \,|\, (X_n, Y_n) = (x_n, y_n)\big) &= \mathbb{P}(Y_{n+1} = y_{n+1} \,|\, Y_n = y_n) \\ &= P(y_n, y_{n+1}).\end{aligned}$$

It should be mentioned that the term "Markovian coupling" is used in [4] to describe what we just defined as faithful coupling. However, since we actually want to focus on couplings that are not faithful (but may still be Markov chains – as both the example in Section 5 and the one in [3] are), we want to make this distinction by adopting the definitions in [5] and deviate from the notions in [4].

In order to understand what makes faithful couplings special, note that in general a coupling of two copies $\mathbf{X}$ and $\mathbf{Y}$ of a Markov chain with transition probabilities $\{P(r, s);\ r, s \in S\}$ fulfills

$$\begin{aligned}\sum_{y_n\in S} \mathbb{P}\big(X_{n+1} = x_{n+1} \,|\, (X_n, Y_n) = (x_n, y_n)\big) \cdot \mathbb{P}(Y_n = y_n \,|\, X_n = x_n) \\ = P(x_n, x_{n+1}),\end{aligned}$$

for all $x_n, x_{n+1} \in S$, $n \in \mathbb{N}_0$, and likewise

$$\begin{aligned}\sum_{x_n\in S} \mathbb{P}\big(Y_{n+1} = y_{n+1} \,|\, (X_n, Y_n) = (x_n, y_n)\big) \cdot \mathbb{P}(X_n = x_n \,|\, Y_n = y_n) \\ = P(y_n, y_{n+1}),\end{aligned}$$

for all $y_n, y_{n+1} \in S$, $n \in \mathbb{N}_0$. So the extra condition on a faithful coupling amounts to $\mathbb{P}\big(X_{n+1} = x_{n+1} \,|\, (X_n, Y_n) = (x_n, y_n)\big)$ being constant in y_n and $\mathbb{P}\big(Y_{n+1} = y_{n+1} \,|\, (X_n, Y_n) = (x_n, y_n)\big)$ being constant in x_n.

It is immediate to check that any faithful coupling can indeed be transformed into a sticky coupling by just letting the chains run according to the given coupling until they meet and then run them together as two *identical* copies of the same chain, without affecting the marginals. Exploiting this fact leads to the estimate

$$\|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}} \leq \mathbb{P}(\tau > n) = 1 - \mathbb{P}(\tau \leq n) \tag{1}$$

for any faithful coupling of two copies, $\mathbf{X}$ started in x and $\mathbf{Y}$ started in y, where

$$\tau := \inf\{n \geq 0;\ X_n = Y_n\}$$

is the *first meeting time* of the coupled chains (cf. Thm. 1 in [5]).

3 Chains that meet and separate

If two Markov chains are coupled, but the coupling is not sticky, clearly they can meet in one state and separate afterwards. If the coupling is not faithful, i.e. violates the conditions laid down in Definition 4, it turns out, sometimes, to be the case that it cannot be transformed into a sticky coupling by simply glueing together the two chains once they meet. As a byproduct, Häggström [3] observed an even more drastic form of incompatibility of two coupled copies of a chain that meet. He gives an example of a finite Markov chain with an intriguing property: Two copies of the chain, started in different states x and y, can be coupled in such a way that they meet almost surely, but their distributions stay apart from one another for all time, more precisely (see Prop. 4.1 in [3]):

Proposition 3.1
There exists a finite state Markov chain such that for two of its states x and y we have that

$$\lim_{n\to\infty} \|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}} > 0,$$

while on the other hand there exists a Markovian coupling of the chains $\mathbf{X} = (X_n)_{n\in\mathbb{N}_0}$ and $\mathbf{Y} = (Y_n)_{n\in\mathbb{N}_0}$, starting at $X_0 = x$ and $Y_0 = y$, with the property that their first meeting time $\tau = \inf\{n \geq 0;\ X_n = Y_n\}$ is finite with probability 1.

Note that for any Markov chain and any two states x and y, the sequence $(\|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}})_{n\in\mathbb{N}_0}$ is non-increasing. This together with the fact that the total variation distance is always non-negative, guarantees the existence of $\lim_{n\to\infty} \|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}}$.

In fact, the Markov chain in the example given in [3] comprises only 6 states and the first meeting time is a.s. less than or equal to 2. The chain features two

absorbing states (that absorb either copy after two steps, no matter if started in x or y) and the probability that the chain started in x ends up in those is $2-\sqrt{2}$ and $\sqrt{2}-1$ respectively. For the chain started in y it is precisely the other way round. So for $n \geq 2$, $P^n(x,.)$ and $P^n(y,.)$ are unchanging and have a total variation distance of $3-2\sqrt{2} \approx 0.17153$. As mentioned in the introduction, we will call Markov chains that have the property described in Proposition 3.1 to be *segregating the states x and y*.

From the convergence theorem we know that such things cannot happen for irreducible finite Markov chains (not even periodic ones). In fact, even if the chain is either reducible or infinite, a coupling that lets two copies, started in different states, meet almost surely, while the total variation distance of their distributions stays bounded away from 0 for all time, cannot be faithful due to the coupling inequality (1).

4 The upper bound

From the previous section we know that there exist finite reducible Markov chains that *segregate two states x and y*. A central question in this respect is how large a total variation distance between $P^n(x,.)$ and $P^n(y,.)$ can be kept, under the condition that two copies started in x and y respectively can be coupled to meet in finite time almost surely. Let us denote this supremum by δ. The example in [3] shows $\delta \geq 3-2\sqrt{2}$; the following theorem establishes $\frac{1}{2}$ as an upper bound on δ.

Theorem 4.1
Consider a Markov chain on the countable state space S and two fixed states x and y. Further, we denote by $\mathbf{X} = (X_n)_{n\in\mathbb{N}_0}$ and $\mathbf{Y} = (Y_n)_{n\in\mathbb{N}_0}$ two coupled copies of the chain, started in x and y respectively, and their first meeting time by τ. If $\mathbf{X}$ and $\mathbf{Y}$ can be coupled in such a way that $\mathbb{P}(\tau < \infty) = 1$, we necessarily have

$$\lim_{n\to\infty} \|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}} \leq \tfrac{1}{2}.$$

This result will immediately be implied by the following proposition:

Proposition 4.2
Consider $\mathbf{X} = (X_n)_{n\in\mathbb{N}_0}$, $\mathbf{Y} = (Y_n)_{n\in\mathbb{N}_0}$ and τ as above. Then it follows that

$$\|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}} \leq 1 - \tfrac{1}{2}\cdot\mathbb{P}(\tau \leq n) \quad \textit{for all } n \in \mathbb{N}_0.$$

Proof: Fix $n \in \mathbb{N}_0$ and a subset $A \subseteq S$ by means of which we define the processes $(M_t)_{t=0}^n$ and $(N_t)_{t=0}^n$ given by

$$M_t := \mathbb{P}\left(X_n \in A \,|\, X_t\right) = P^{n-t}(X_t, A),$$

and

$$N_t := \mathbb{P}\left(Y_n \in A \,|\, Y_t\right) = P^{n-t}(Y_t, A).$$

It is easily checked that these processes are martingales (with respect to the filtrations generated by $\mathbf{X}$ and $\mathbf{Y}$ respectively). We further let B_x and B_y denote the events that $M_t > \frac{1}{2}$ and $N_t < \frac{1}{2}$ respectively for all $0 \le t \le n$. As $X_t = Y_t$ implies $M_t = N_t$, it follows that the event $\{\tau \le n\}$ is contained in the union of B_x^c and B_y^c.

Next, we define

$$\tau_x := \inf\{0 \le t \le n;\ M_t \le \tfrac{1}{2}\} \quad \text{and} \quad \tau_y := \inf\{0 \le t \le n;\ N_t \ge \tfrac{1}{2}\},$$

where the infimum is understood to be n if the corresponding set is empty. Note that τ_x and τ_y are stopping times for M_t and N_t respectively. Since $(M_t)_{t=0}^n$ and $(N_t)_{t=0}^n$ are bounded martingales, the Optional Stopping Theorem (see for example Cor. 17.7 in [4]) gives the estimates

$$P^n(x,A) = \mathbb{E}\, M_0 = \mathbb{E}\, M_{\tau_x} \le \tfrac{1}{2} \cdot \mathbb{P}(B_x^c) + 1 \cdot \mathbb{P}(B_x) = 1 - \tfrac{1}{2} \cdot \mathbb{P}(B_x^c)$$

$$\text{and} \quad P^n(y,A) = \mathbb{E}\, N_0 = \mathbb{E}\, N_{\tau_y} \ge \tfrac{1}{2} \cdot \mathbb{P}(B_y^c).$$

Combining these two inequalities, we get

$$P^n(x,A) - P^n(y,A) \le 1 - \tfrac{1}{2}\left(\mathbb{P}(B_x^c) + \mathbb{P}(B_y^c)\right) \le 1 - \tfrac{1}{2} \cdot \mathbb{P}(\tau \le n).$$

Finally, maximizing the left-hand side over all subsets $A \subseteq S$ yields

$$\|P^n(x,\cdot) - P^n(y,\cdot)\|_{TV} \le 1 - \tfrac{1}{2} \cdot \mathbb{P}(\tau \le n),$$

as claimed. □

Proof of Theorem 4.1: If $\mathbf{X}$ and $\mathbf{Y}$ are now coupled in such a way that $\mathbb{P}(\tau < \infty) = 1$, we can immediately conclude from Proposition 4.2 that

$$\lim_{n\to\infty} \|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}} \le \lim_{n\to\infty} 1 - \tfrac{1}{2} \cdot \mathbb{P}(\tau \le n) = \tfrac{1}{2}.$$ □

Remark 1

Reading carefully through the proof of Proposition 4.2, one may notice that the martingale argument does not use our general assumptions, namely time-homogeneity and discrete time. To transfer the result to a Markov chain in continuous time that is not necessarily time-homogeneous, we can simply replace $\mathbb{N}_0$ by $\mathbb{R}_{\ge 0}$ and $\{0,\dots,n\}$ by $[0,n]$ in both statement and proof. There is however a small technical subtlety caused by the infimum in the definition of the stopping times.

Given transition probabilities that are *standard* (cf. the beginning of Chap. 5 in the book of Freeman [2]), meaning that

$$\lim_{\varepsilon \downarrow 0} \mathbb{P}(X_{t+\varepsilon} = j \,|\, X_t = i) = \delta_{ij}, \quad \text{for all } i,j \in S,\ t \ge 0,$$

the functions

$$t \mapsto \mathbb{P}(X_t = j \,|\, X_s = i),\ t \ge s$$

become right continuous for all $i, j \in S$, $s \geq 0$ (cf. Lemma (9) in [2]), which in turn implies continuity from the right of the martingales $(M_t)_{t\in[0,n]}$ and $(N_t)_{t\in[0,n]}$ as defined above. This lets our argument go through and establishes Proposition 4.2 and therefore also Theorem 4.1 for this more general setting.

Besides this generalization, the statement from Proposition 4.2 can also be used to get upper bounds on mixing times – similar to the usual approach, see for example Cor. 5.3 in [4] – replacing the coupling inequality (1) as starting point. In doing so, we pay by an additional factor $\frac{1}{2}$ in front of $\mathbb{P}(\tau \leq n)$, but can in return employ any kind of coupling, not only faithful ones. It remains, however, to be seen whether this will ever turn out useful in practice as basically all standard coupling constructions are faithful.

Proposition 4.3
Consider a Markov chain $\mathbf{X} = (X_n)_{n\in\mathbb{N}_0}$ with the property that the distribution of X_n converges to a fixed distribution π irrespectively of the starting distribution. Further, suppose that for some $\alpha \in (0, 1]$ and each pair of states $x, y \in S$ there exists a (not necessarily faithful or even Markovian) coupling $((X_n, Y_n))_{n\in\mathbb{N}_0}$ of two copies of the chain started in x and y respectively, such that the first meeting time τ of the two coupled processes fulfills $\mathbb{P}(\tau \leq n) \geq \alpha$. Then

$$t_{\text{mix}} \leq n \cdot \left\lceil \frac{\log(\frac{1}{4})}{\log(1 - \frac{\alpha}{2})} \right\rceil .$$

Proof: From Proposition 4.2 we can conclude that

$$\overline{d}(n) = \sup_{x,y\in S} \|P^n(x, .) - P^n(y, .)\|_{\text{TV}} \leq 1 - \tfrac{\alpha}{2}.$$

Consequently, as $\overline{d}$ is submultiplicative and dominates d, we get for any $k \in \mathbb{N}$:

$$d(kn) \leq \overline{d}(kn) \leq \overline{d}(n)^k \leq (1 - \tfrac{\alpha}{2})^k.$$

Choosing $k \geq \frac{\log\left(\frac{1}{4}\right)}{\log\left(1-\frac{\alpha}{2}\right)}$, this estimate becomes

$$d(kn) \leq (1 - \tfrac{\alpha}{2})^{\log\left(\frac{1}{4}\right)/\log\left(1-\frac{\alpha}{2}\right)} = \frac{1}{4}.$$

□

5 A simple example that narrows the gap

In this section, we will present a readily accessible finite state Markov chain that improves on the value of $3 - 2\sqrt{2}$ established in [3]. First, let us prepare an auxiliary result, which will be needed to assess the total variation distance in this example of a finite segregating Markov chain.

Consider a sequence of independent Bernoulli trials, each with success probability $p < 1$. The distribution of the number of successful attempts until r failures have occurred is called the negative binomial distribution with parameters r and p and commonly denoted by $\text{NB}(r, p)$.

Lemma 5.1
With $\mu := \mathrm{NB}(1,p)$ and $\nu := \mathrm{NB}(2,p)$, we claim that

$$\|\mu-\nu\|_{\mathrm{TV}} = \sum_{k=0}^{\lfloor \frac{p}{1-p} \rfloor} \big(\mu(k)-\nu(k)\big) = \big(\lfloor \tfrac{p}{1-p} \rfloor + 1\big)\cdot(1-p)\cdot p^{\lfloor \frac{p}{1-p} \rfloor + 1} \qquad (2)$$

for any $p \in [0,1)$. If $p = \frac{m}{m+1}$, with $m \in \mathbb{N}_0$, this value simplifies to $p^{\frac{1}{1-p}}$.

Proof: A standard calculation shows that for two probability distributions μ and ν on a discrete space S, their total variation distance can be calculated as

$$\|\mu-\nu\|_{\mathrm{TV}} = \sum_{\substack{x\in S \\ \mu(x)\geq\nu(x)}} \big(\mu(x)-\nu(x)\big),$$

see for example Remark 4.3 in [4]. For $\mu = \mathrm{NB}(1,p)$ and $\nu = \mathrm{NB}(2,p)$, we have:

$$\mu(k) = p^k\,(1-p) \quad \text{and} \quad \nu(k) = (k+1)\,p^k\,(1-p)^2, \quad \text{for all } k \in \mathbb{N}_0,$$

which implies $\mu(k) \geq \nu(k)$ for $k \leq \lfloor \frac{p}{1-p} \rfloor =: m$.

Consequently, using the well-known formulas for a finite geometric sum and its derivative, we can compute the total variation distance and get

$$\begin{aligned}\|\mu-\nu\|_{\mathrm{TV}} &= \sum_{k=0}^{m} p^k\,(1-p) - \sum_{k=0}^{m} (k+1)\,p^k\,(1-p)^2 \\ &= 1 - p^{m+1} - \big[1 - p^{m+2} - (m+2)\,p^{m+1}\,(1-p)\big] \\ &= (m+1)\,(1-p)\,p^{m+1}.\end{aligned}$$

If $p = \frac{m}{m+1}$, for some $m \in \mathbb{N}_0$, the number $\frac{p}{1-p} = m$ is integer and an elementary simplification of the expression to the right in (2) verifies the final claim. □

Let us now use this lemma to establish the following result:

Proposition 5.2
For all $\varepsilon > 0$, there exists a Markov chain segregating two states x and y such that

$$\lim_{n\to\infty} \|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}} > \frac{1}{\mathrm{e}} - \varepsilon.$$

Proof: Let us consider the finite reducible MC depicted in Figure 1. The state space S comprises $3m+5$ states, among which the two initial states x and y as well as the $m+2$ absorbing states labeled $0,\ldots,m$ and $>$.

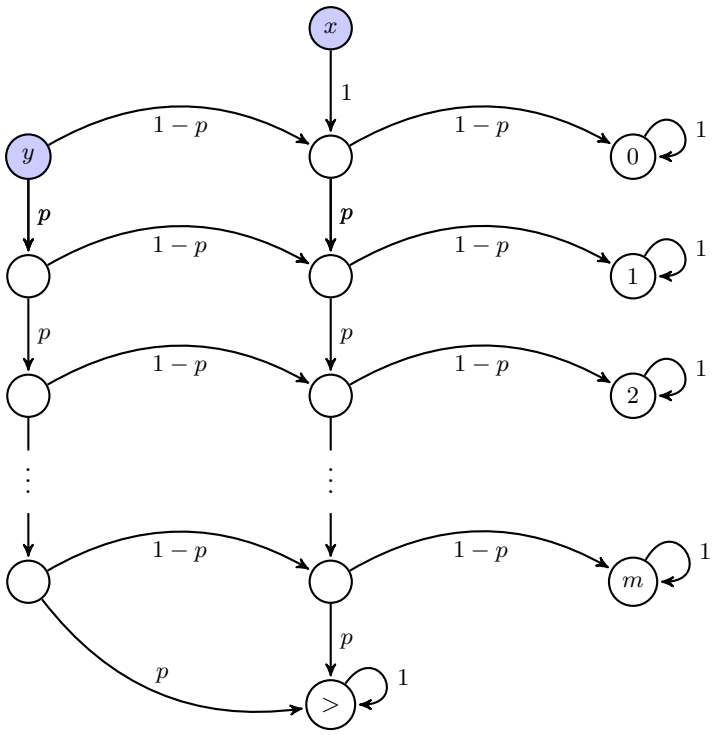

Figure 1: A segregating MC allowing for a fairly large total variation distance of two copies started in x and y respectively.

It is immediate to check that the two copies, started in x respectively y, will hit an absorbing state after at most $m+2$ steps and that on $\{0,\dots,m\}$ the distributions $P^{m+2}(x,.)$ and $P^{m+2}(y,.)$ coincide with $\mathrm{NB}(1,p)$ and $\mathrm{NB}(2,p)$ respectively. The probabilities to end up in the state labeled $>$ are

$$P^{m+2}(x,>) = \mathbb{P}(Z_1 > m) \quad \text{and} \quad P^{m+2}(y,>) = \mathbb{P}(Z_2 > m),$$

where $Z_1 \sim \mathrm{NB}(1,p)$ and $Z_2 \sim \mathrm{NB}(2,p)$.

Choosing $p = \frac{m}{m+1}$, we can conclude from Lemma 5.1 that

$$\|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}} = p^{\frac{1}{1-p}},$$

for all $n \geq m+2$. Writing $q = \frac{1}{1-p}$ shows

$$\lim_{p\to 1} p^{\frac{1}{1-p}} = \lim_{q\to\infty} \left(1 - \tfrac{1}{q}\right)^q = \frac{1}{\mathrm{e}}$$

and hence taking m large enough will establish the claim if we can present a coupling that ensures that the two copies started in x and y will meet with probability 1, either before or when they hit an absorbing state.

In order to establish such a coupling, let $\mathbf{Y} = (Y_n)_{n\in\mathbb{N}_0}$ be a copy of the MC started in y. The copy $\mathbf{X} = (X_n)_{n\in\mathbb{N}_0}$ started in x mimicks all movements of $\mathbf{Y}$, with the delay of one step, until it finally hits an absorbing state: First, it will move downwards until the two processes meet – in particular this implies that its first step is downwards with probability 1, as $x \neq y$. Then, once $X_n = Y_n$ for some $1 \leq n \leq m+1$, the next step of the process $\mathbf{X}$ is to move to the right to an absorbing state, i.e. $X_{n+1} = n-1$. If $\mathbf{Y}$ never moves to the right, neither does $\mathbf{X}$ and both end up in the state $>$.

First of all, we need to check whether the two coordinate processes are indeed copies of the MC given in Figure 1: It is obvious from our construction that it suffices to verify this for the process $\mathbf{X}$. The way $\mathbf{X}$ is defined – to move downwards in the first step and then always imitate the previous move of $\mathbf{Y}$ until ending up in an absorbing state – gives the right marginals due to the structure of the MC: As all the non-absorbing states apart from x have the same transition probabilities (p downwards and $1-p$ to the right), $\mathbf{X}$ performs indeed a random walk on the graph in Figure 1 according to the transition probabilites of the MC. Note that there is just this one way $\mathbf{Y}$ can end up in an absorbing state before it meets $\mathbf{X}$, namely if it moves downwards only. Then, however, $\mathbf{X}$ copies this behavior and ends up in the state $>$ as well, so the coupling guarantees $\tau \leq m+2$ with probability 1. This trivially implies the almost sure finiteness of the first meeting time τ and in conclusion the claim that the MC segregates the two states x and y. □

To grasp the importance of faithfulness in this context, it is worth noting that the coupling in our example – likewise the one in [3] – is in fact Markovian but not faithful.

6 Closing the gap

In this last section, we want to improve further on the example from the previous section in order to determine the true value of the universal constant δ. These efforts amount to the following:

Theorem 6.1
The value of δ, denoting the supremum of $\lim_{n\to\infty}\|P^n(x,.) - P^n(y,.)\|_{\mathrm{TV}}$ taken over all segregating Markov chains and segregated states x and y, is $\frac{1}{2}$.

In view of Theorem 4.1, the final step to derive this result lies in proving that for any $\varepsilon > 0$, a value of at least $\frac{1}{2} - \varepsilon$ can actually be attained. In order to do so, we will focus on (reducible) Markov chains with a special structure which allows us to consider even simpler chains in finite time instead: When talking about a Markov chain with finite time horizon $\mathbf{X} = (X_t)_{t=0}^T$ on state space S with transition probabilities $\{P(r,s);\ r,s \in S\}$, we are in fact thinking of a different chain, namely the reducible Markov chain $\mathbf{Y} = (Y_n)_{n\in\mathbb{N}_0}$ on the state space $S \times \{0,\dots,T\}$, with transition probabilities

$$\mathbb{P}(Y_{n+1} = (s, n+1) \,|\, Y_n = (r,n)) = P(r,s) \quad \text{for all } r,s \in S,\ n \in \{0,\dots,T-1\}$$

$$\text{and} \quad \mathbb{P}(Y_{n+1} = (s,T) \,|\, Y_n = (s,T)) = 1.$$

In other words, we consider the evolution in time as new generations of states and stop the original chain at time T by making all states corresponding to time T absorbing. Incidentally, the example given in [3] is also of this kind; it corresponds to a two state Markov chain with $T = 2$.

So for the remainder of this article, we actually consider Markov chains in discrete, finite time and with finite state space only. With this simplification in mind, we want to prove the following:

Theorem 6.2
For any $\varepsilon > 0$ there exists a Markov chain $\mathbf{X}$, two states x and y and a positive integer T such that

$$\|P^T(x,.) - P^T(y,.)\|_{\mathrm{TV}} \geq \frac{1}{2} - \varepsilon$$

and such that there is a coupling of two copies of the chain, $((X_t, Y_t))_{t=0}^T$, with initial states x and y respectively satisfying

$$\mathbb{P}\left(X_t = Y_t \text{ for some } 0 \leq t \leq T\right) = 1.$$

Before proceeding with a proof of this theorem, we want to introduce an alternative way to view segregating couplings. Let $\mathbf{X}$ be any discrete-time Markov chain with a finite state space S, and fix two states $x, y \in S$ and a positive integer T.

For any coupling of two copies of this chain, $((X_t, Y_t))_{n=0}^T$, started in $X_0 = x$ and $Y_0 = y$ respectively, one can consider the corresponding *meeting probability*

$$\mathbb{P}\left(X_t = Y_t \text{ for some } 0 \leq t \leq T\right). \tag{3}$$

A natural question to ask in this setting is how large one can make this probability for a given (finite) Markov chain and given x, y and T by maximizing over all such couplings. Let $\mathcal{X}$ denote the subset of $\{x\} \times S^T$ consisting of all possible trajectories of $\mathbf{X} = (X_t)_{t=0}^T$, and respectively $\mathcal{Y} \subseteq \{y\} \times S^T$ for $\mathbf{Y} = (Y_t)_{t=0}^T$. Observe that any coupling of $\mathbf{X}$ and $\mathbf{Y}$ is determined by the values $(p_{\mathbf{xy}})_{\mathbf{x}\in\mathcal{X},\mathbf{y}\in\mathcal{Y}}$, where $p_{\mathbf{xy}}$ denotes the probability of the event that both $\mathbf{X} = \mathbf{x} = (x_t)_{t=0}^T$ and $\mathbf{Y} = \mathbf{y}$. We will denote the maximal value of (3) by $\mathcal{C}_T(x,y)$ and call it the *optimal meeting probability*.

While finding explicit couplings that maximize the meeting probability can quickly become cumbersome as the number of possible trajectories grows, it turns out that the problem of optimizing the meeting probability has a useful dual, which allows us to determine $\mathcal{C}_T(x,y)$ without having to deal with the couplings directly. This duality corresponds to the idea of max-flow/min-cut and König's theorem in combinatorial optimization.

Definition 5
Let $\mathcal{A} = (A_t)_{t=0}^T$ be a sequence of subsets of S, the (finite) state space of the considered Markov chain $\mathbf{X}$. We will refer to any such sequence as a *separating*

sequence. We define the *separation* of any separating sequence as

$$\begin{aligned}\mathcal{S}_T^{\mathcal{A}}(x,y) = \mathbb{P}\left(X_t \in A_t \text{ for all } 0 \le t \le T \,|\, X_0 = x\right) \\ + \mathbb{P}\left(X_t \notin A_t \text{ for all } 0 \le t \le T \,|\, X_0 = y\right).\end{aligned} \tag{4}$$

We say that the separating sequence is *non-trivial* if both summands on the right-hand side in (4) are non-zero. We further define the *optimal separation* $\mathcal{S}_T(x,y)$ as the maximum separation over all possible separating sequences.

It is not too hard to see that the optimal meeting probability and optimal separation are related. Specifically, for any coupling $((X_t,Y_t))_{t=0}^T$ such that $(X_0,Y_0)=(x,y)$ and any separating sequence $\mathcal{A}=(A_t)_{t=0}^T$ we have

$$\begin{aligned}&\mathbb{P}\left(X_t = Y_t \text{ for some } 0 \le t \le T\right) \\ &\quad\le \mathbb{P}\left(X_t \notin A_t \text{ or } Y_t \in A_t \text{ for some } 0 \le t \le T\right) \\ &\quad\le \mathbb{P}\left(X_t \notin A_t \text{ for some } 0 \le t \le T\right) + \mathbb{P}\left(Y_t \in A_t \text{ for some } 0 \le t \le T\right) \\ &\quad= 2 - \mathcal{S}_T^{\mathcal{A}}(x,y).\end{aligned}$$

Maximizing the meeting probability over all possible couplings of two copies, started in x,y respectively, and minimizing the upper bound by maximizing the separation $\mathcal{S}_T^{\mathcal{A}}(x,y)$ over all separating sequences yields

$$\mathcal{C}_T(x,y) \le 2 - \mathcal{S}_T(x,y). \tag{5}$$

However, as the following theorem shows, the connection between optimal meeting probability and optimal separation goes beyond this inequality.

Theorem 6.3
Given an arbitrary but fixed finite Markov chain $\mathbf{X}$*, two states* $x,y \in S$ *and time horizon* T*, we have*

$$\mathcal{C}_T(x,y) = 2 - \mathcal{S}_T(x,y). \tag{6}$$

Proof: A simple way to prove the reverse inequality is to employ the max-flow min-cut theorem, in the same way it can be used to prove Strassen's monotone coupling theorem. Starting from the sets $\mathcal{X}$ and $\mathcal{Y}$ as above, we build the following directed graph, which will be denoted by $\vec{G}=(V,\vec{E})$:

First we let each $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$ be represented by a node. Then we add two further nodes: a source s and a sink t. When it comes to the directed edges, there will be an arrow $(s,\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$ and $(\mathbf{y},t)$ for all $\mathbf{y} \in \mathcal{Y}$. Additionally, we include the edge $(\mathbf{x},\mathbf{y})$, if the two trajectories $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$ share at least one state, i.e. $x_t = y_t$ for some $0 \le t \le T$; in the sequel, we will write this as $\mathbf{x} \sim \mathbf{y}$.

Finally, we have to assign capacities to these directed edges: The edges $(s,\mathbf{x})$ and $(\mathbf{y},t)$ will get capacities $\mathbb{P}(\mathbf{X}=\mathbf{x})$ and $\mathbb{P}(\mathbf{Y}=\mathbf{y})$ respectively. All edges going in between $\mathcal{X}$ and $\mathcal{Y}$ get capacity 1, see Figure 2 below for an illustration.

Let us now consider the minimum cut problem on $\vec{G}$. From the fact that the cut $(\{s\}, V\setminus\{s\})$ has value 1, we know that we can focus on cuts that are

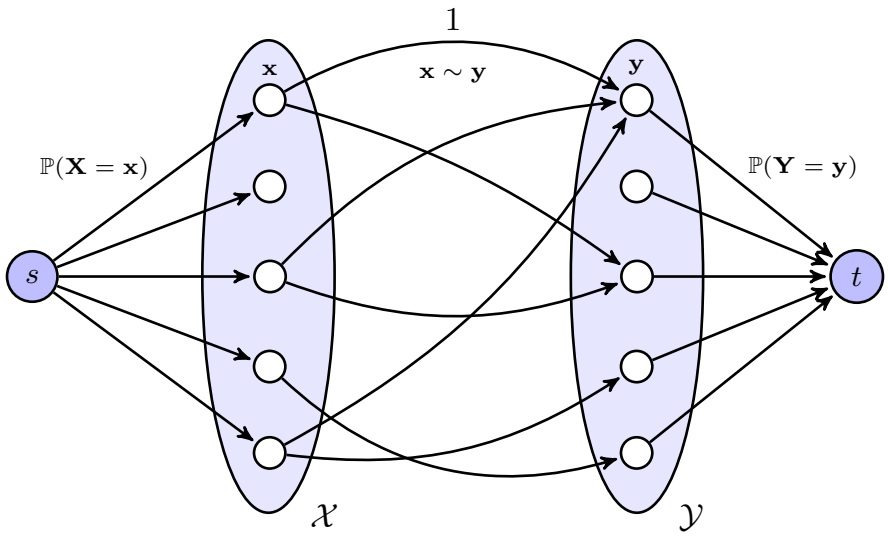

Figure 2: The auxiliary graph $\vec{G}$ to which we apply the max-flow min-cut theorem.

not cutting any edges going in between $\mathcal{X}$ and $\mathcal{Y}$ when trying to find a minimal one. For a cut (B, B^{c}) of this kind, with $s \in B$ and $t \in B^{\mathrm{c}}$ say, we note that $\mathbf{x} \sim \mathbf{y}$ can not occur for $\mathbf{x} \in \mathcal{X} \cap B$ and $\mathbf{y} \in \mathcal{Y} \cap B^{\mathrm{c}}$, due to the assumption that no such edges are cut. Furthermore, the edges cut incident to s have at least the value $\mathbb{P}(X \in \mathcal{X} \cap B^{\mathrm{c}})$, the ones incident to t at least $\mathbb{P}(Y \in \mathcal{Y} \cap B)$.

Let us define the set sequence $\mathcal{A} = (A_t)_{t=0}^{T}$ by

$$A_t := \{x_t;\ \mathbf{x} = (x_s)_{s=0}^{T} \in \mathcal{X} \cap B\},\ \text{for all } 0 \leq t \leq T,$$

in order to bound the value of the given cut from below by

$$\begin{aligned}
&\mathbb{P}(X \in \mathcal{X} \cap B^{\mathrm{c}}) + \mathbb{P}(Y \in \mathcal{Y} \cap B) \\
&\qquad \geq 1 - \mathbb{P}(X_t \in A_t \text{ for all } 0 \leq t \leq T) + 1 - \mathbb{P}(Y_t \notin A_t \text{ for all } 0 \leq t \leq T) \\
&\qquad = 2 - \mathcal{S}_T^{\mathcal{A}}(x, y) \\
&\qquad \geq 2 - \mathcal{S}_T(x, y).
\end{aligned}$$

Consequently, as this bound applies to any minimal cut, using the max-flow min-cut theorem (see for example Thm. 1, Chap. III in [1]) we are guaranteed the existence of a maximal flow of value at least $2 - \mathcal{S}_T(x, y)$. Let us denote the respective flow through the edge corresponding to $\mathbf{x} \sim \mathbf{y}$ by $q_{\mathbf{xy}}$.

We can use this maximal flow to establish a coupling of $\mathbf{X}$ and $\mathbf{Y}$ in the same vein as in Doeblin's coupling lemma (see for instance Prop. 4.7 in [4]): First we let $\mathbf{X} = \mathbf{x}$ and $\mathbf{Y} = \mathbf{y}$ simultaneously with probability $q_{\mathbf{xy}}$ for all $\mathbf{x} \sim \mathbf{y}$. Then

we define X to follow the trajectory $\mathbf{x}$ with the remaining probability

$$\mathbb{P}(\mathbf{X} = \mathbf{x}) - \sum_{\mathbf{y}:\ \mathbf{x} \sim \mathbf{y}} q_{\mathbf{xy}}$$

and similarly Y to follow the trajectory $\mathbf{y}$ with probability

$$\mathbb{P}(\mathbf{Y} = \mathbf{y}) - \sum_{\mathbf{x}:\ \mathbf{x} \sim \mathbf{y}} q_{\mathbf{xy}},$$

independently, for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$.

From the flow constraints, we know that all these probabilities are in $[0, 1]$ and the resulting coupling satisfies $\mathbb{P}(\mathbf{X} \sim \mathbf{Y}) \geq 2 - \mathcal{S}_T(x, y)$. The theorem then follows by combining this inequality with (5). □

One can observe that, as the left-hand side of (6) is a probability and hence at most 1, we must always have optimal separation at least 1. Indeed, we can obtain separation equal to 1, for instance by taking $A_t = S$ for all $0 \leq t \leq T$. Recall that a separating sequence is called non-trivial if the probabilities that $X_t \in A_t$ for all $0 \leq t \leq T$ given $X_0 = x$ and $X_t \notin A_t$ for all $0 \leq t \leq T$ given $X_0 = y$ are both non-zero. Clearly, any trivial separating sequence has separation at most 1, so it follows from Theorem 6.3 that for any finite state Markov chain in discrete time, any x, y and T as above, the following two statements are equivalent:

(a) The meeting probability under optimal coupling of two copies, started in x, y respectively, is 1.

(b) For all non-trivial separating sequences $\mathcal{A} = (A_t)_{t=0}^T$, we get $\mathcal{S}_T^{\mathcal{A}}(x, y) \leq 1$.

In order to get acquainted with the idea behind the concept of separation, let us take a look at the simplest non-trivial example:

Example 6.1
Let $0 < \alpha \leq \frac{1}{2}$. Consider the Markov chain $\mathbf{X}$ with state space $\{0, 1\}$ and transition probabilities $P(0, 1) = P(1, 0) = \alpha$ and $P(0, 0) = P(1, 1) = 1 - \alpha$, and take $x = 0,\ y = 1$. As mentioned above, the case where $T = 2$ is Häggström's [3] example of a segregating Markov chain.

As this chain only has two states, any non-trivial separating sequence must have $A_t = \{0\}$ or $\{1\}$ for each $0 \leq t \leq T$. As $\alpha \leq \frac{1}{2}$, the non-trivial separating sequence given by $A_t = \{0\}$ for all t is obviously best possible. It is immediate to check that its separation equals $2(1 - \alpha)^T$. Hence, the optimal meeting probability of two copies started in states 0 and 1 respectively is 1 if and only if $2(1 - \alpha)^T \leq 1$. Using induction, one can easily check that for this chain we have

$$\|P^T(0, .) - P^T(1, .)\|_{\mathrm{TV}} = (1 - 2\alpha)^T.$$

So by choosing $\alpha = \alpha(T)$ such that $2(1 - \alpha)^T = 1$, we obtain a Markov chain that segregates the states 0 and 1 with total variation distance $(2^{1-1/T} - 1)^T$, which tends to $\frac{1}{4}$ as $T \to \infty$.

The next example is supposed to illustrate that reducible Markov chains obtained from irreducible and aperiodic finite chains in finite time, in the way described before Theorem 6.2, usually do segregate any two states:

Example 6.2
Let $\mathbf{X}$ be any irreducible aperiodic Markov chain with a finite state space S, and let x and y be any two states. Pick $n \in \mathbb{N}$ and $\varepsilon > 0$ such that $P^n(x', y') \geq \varepsilon$ for all $x', y' \in S$. Then, for any non-trivial separating sequence $(A_t)_{t=0}^{nk}$, we have

$$\mathcal{S}_{nk}^{\mathcal{A}}(x, y) \leq 2\,(1 - \varepsilon)^k$$

for any $k \in \mathbb{N}$. Hence, by taking $T = nk$ for a sufficiently large k it follows that the optimal meeting probability during $[0, T]$ is 1. This shows that, unless $\|P^T(x, .) - P^T(y, .)\|_{\mathrm{TV}} = 0$, the Markov chain segregates the two states x and y, choosing T sufficiently large.

We now turn to the proof of Theorem 6.2. Let $\mathbf{X}$ be the Markov chain with state space $\{0, 1, \ldots, L\}$ for some positive integer L and transition probabilities $P(0, 1) = P(L, L-1) = 1 - P(0, 0) = 1 - P(L, L) = \alpha$ and $P(i, i+1) = P(i, i-1) = \frac{1}{2}$ for all $0 < i < L$, see Figure 3. Such chains – with $S = \{0, 1, \ldots, L\}$ and the special property that $|X_{t+1} - X_t| \leq 1$ a.s. for all t – are commonly called birth-and-death chains, cf. Section 2.5 in [4]. We shall here consider the optimal meeting probability of this chain given the starting states $x = 0$ and $y = L$. We will show that we can obtain segregation between 0 and L with total variation arbitrarily close to $\frac{1}{2}$ by choosing L, T and α appropriately.

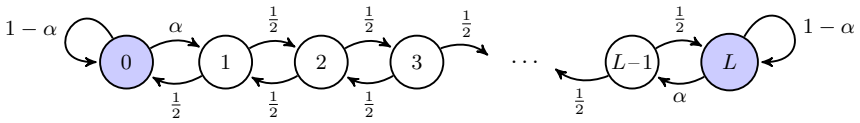

Figure 3: A birth-and-death chain on $\{0, 1, \ldots, L\}$ as basic building block.

The qualitative behavior of this chain for small α is easy to describe: Most of the time, the process is either at 0 or L. Occasionally, that is at rate α, the process takes one step inwards and typically moves around for order L steps before hitting one of the marginal states 0 or L, more precisely the expected time to reach $\{0, L\}$, when starting in state 1 equals $L - 1$. With probability $1 - \frac{1}{L}$ it will return to the same side it started at, with probability $\frac{1}{L}$ it will cross over to the opposite side (check the analysis of the so-called gambler's ruin problem in Section 2.1 in [4] for the explicit calculations).

In preparation to the in-depth analysis of the separation of the states 0 and L, let us collect some general estimations for this chain in the following lemma that will come in useful later on.

Lemma 6.4
(a) For any $1 \leq i \leq L - 1$ and $t \geq 0$ we have $P^t(0, i) = P^t(L, L - i) \leq 2\alpha$.

(b) For any $t \geq 0$ we have

$$P^t(0,0) = P^t(L,L) = \frac{1}{2} + \frac{1}{2}\left(1 - \frac{2\alpha}{L}\right)^t + O(L\alpha), \tag{7}$$

$$\|P^t(0,.) - P^t(L,.)\|_{\mathrm{TV}} = \left(1 - \frac{2\alpha}{L}\right)^t + O(L\alpha), \tag{8}$$

(c) For any $0 \leq k \leq L-1$ and $t \geq 0$ we have

$$\begin{aligned} &\mathbb{P}\left(X_{t'} \leq k \text{ for all } 0 \leq t' \leq t \,|\, X_0 = 0\right) \\ &\quad = \mathbb{P}\left(X_{t'} \geq L-k \text{ for all } 0 \leq t' \leq t \,|\, X_0 = L\right) \\ &\quad = \left(1 - \frac{\alpha}{k+1}\right)^t + O(k\alpha). \end{aligned} \tag{9}$$

Proof:

(a) The first statement easily follows by induction on t using the recursion

$$P^{t+1}(0,i) = P^t(0,i-1)\,P(i-1,i) + P^t(0,i+1)\,P(i+1,i),$$

which holds for any $t \geq 0$ and any $1 \leq i \leq L-1$.

(b) To show the second claim, consider the sequence $a_t = \mathbb{E}\left[X_t \,|\, X_0 = 0\right]$. Using (a), we know that

$$a_t = L \cdot P^t(0,L) + O(L^2\alpha) = L - L \cdot P^t(0,0) + O(L^2\alpha), \tag{10}$$

where the error terms are bounded by $2\,L^2\alpha$ in absolute value. Furthermore, as $\mathbb{E}\left[X_{t+1} - X_t \,|\, X_t\right] = \alpha$ if $X_t = 0$, $-\alpha$ if $X_t = L$ and 0 otherwise, we can infer from (10)

$$\begin{aligned} a_{t+1} &= a_t + \alpha \cdot P^t(0,0) - \alpha \cdot P^t(0,L) \\ &= a_t + \alpha\left(1 - \frac{a_t}{L}\right) - \alpha\,\frac{a_t}{L} + O(L\alpha^2), \end{aligned}$$

which implies

$$a_{t+1} - \frac{L}{2} = \left(1 - \frac{2\alpha}{L}\right) \cdot \left(a_t - \frac{L}{2}\right) + O(L\alpha^2),$$

where the error term is bounded by $4\,L\alpha^2$ in absolute value irrespectively of t. Solving this recursion, using $a_0 = 0$ and $\sum_{k=0}^{t-1}(1 - \frac{2\alpha}{L})^k \leq \frac{L}{2\alpha}$, yields

$$a_t = \frac{L}{2} - \frac{L}{2}\left(1 - \frac{2\alpha}{L}\right)^t + O(L^2\alpha).$$

The estimate (7) immediately follows from (10), which together with part (a) implies (8).

(c) The case $k = 0$ is obvious, so we may assume $k > 0$. Let τ denote the first time $t \geq 0$ for which $X_t = k+1$. Consider the sequence $b_t = \mathbb{E}[X_{t\wedge\tau} \,|\, X_0 = 0]$, where $t \wedge \tau$ denotes the minimum of t and τ. Note that part (a) implies $\mathbb{P}(X_{t\wedge\tau} = i \,|\, X_0 = 0) \leq \mathbb{P}(X_t = i \,|\, X_0 = 0) \leq 2\alpha$ for any $i \in \{1, \ldots, k\}$ and further

$$\begin{aligned} b_t &= (k+1) \cdot \mathbb{P}(X_{t\wedge\tau} = k+1 \,|\, X_0 = 0) + O(k^2\alpha) \\ &= k + 1 - (k+1) \cdot \mathbb{P}(X_{t\wedge\tau} = 0 \,|\, X_0 = 0) + O(k^2\alpha). \end{aligned} \tag{11}$$

With the same reasoning as in part (b), we end up in a similar situation with $b_0 = 0$ and b_t satisfying the recursive formula

$$\begin{aligned} b_{t+1} &= b_t + \alpha \cdot \mathbb{P}(X_{t\wedge\tau} = 0 \,|\, X_0 = 0) \\ &= b_t + \alpha\left(1 - \frac{b_t}{k+1}\right) + O(k\alpha^2). \end{aligned}$$

Solving it gives $b_t = k + 1 - (k+1)(1 - \frac{\alpha}{k+1})^t + O(k^2\alpha)$ and plugging this into (11) completes the proof of part (c), using

$$\mathbb{P}(1 \leq X_{t\wedge\tau} \leq k \,|\, X_0 = 0) = O(k\alpha)$$

and noting that $X_{t\wedge\tau} \leq k$ if and only if $X_{t'} \leq k$ for all $0 \leq t' \leq t$. □

Let us now take a closer look on the optimal separation of states 0 and L in this chain. To make our lives easier, we establish three auxiliary results showing that among the non-trivial separating sequences, there are very simple ones which are essentially best possible as T grows large.

Proposition 6.5
Let L be fixed, and let $\alpha = \alpha(T) = \Theta\left(\frac{1}{T}\right)$. Then, for sufficiently large T, any non-trivial separating sequence $\mathcal{A} = (A_t)_{t=0}^T$ such that $\mathcal{S}_T^{\mathcal{A}}(0, L) > 1$ (if such exist) must satisfy $0 \in A_t$ and $L \notin A_t$ for all $0 \leq t \leq T$.

Proof: By (7), we know that $P^t(0,0) = P^t(L,L) > 1/2$ for any $0 \leq t \leq T$ given T sufficiently large. Hence if $0 \notin A_{t_1}$ and $L \in A_{t_2}$ for some $0 \leq t_1, t_2 \leq L$, then

$$\mathcal{S}_T^{\mathcal{A}}(0, L) \leq 1 - P^{t_1}(0,0) + 1 - P^{t_2}(L,L) < 1.$$

So for T sufficiently large, either $0 \in A_t$ for all $0 \leq t \leq T$ or $L \notin A_t$ for all $0 \leq t \leq T$.

By symmetry, we can assume without loss of generality, that $0 \in A_t$ for all $0 \leq t \leq T$. Note that if there is some t_1 such that $L \in A_{t_1}$, then by part (a) of Lemma 6.4, we get

$$\mathbb{P}(X_t \notin A_t \text{ for all } 0 \leq t \leq T \,|\, X_0 = L) \leq \mathbb{P}(0 < X_{t_1} < L \,|\, X_0 = L) \leq 2L\alpha.$$

So the proposition follows if we can show that there exists a constant $\varepsilon > 0$ such that for T sufficiently large any non-trivial separating sequence satisfies

$$\mathbb{P}(X_t \in A_t \text{ for all } 0 \leq t \leq T \,|\, X_0 = 0) \leq 1 - \varepsilon.$$

First note that since $\mathcal{A}$ is non-trivial, there exists a trajectory $\mathbf{y} \in \{L\} \times S^T$ such that $y_t \notin A_t$ for all $0 \le t \le T$ and $\mathbb{P}(\mathbf{X} = \mathbf{y} \,|\, X_0 = L) > 0$. Further, recall that the chain can only attain $\mathbf{y}$ with positive probability if $|y_{t+1} - y_t| \le 1$ for all $0 \le t \le T-1$.

Next, let $\mathbf{X} = (X_t)_{t=0}^{T+1}$ be a copy of the Markov chain started in $X_0 = 0$. We define the process $\mathbf{X}' = (X'_t)_{t=0}^{T}$ as $X'_t = X_{t+1}$ for all $0 \le t \le T$ if $X_1 = 0$, and otherwise put $X'_0 = 0$ and let this process evolve independently of $\mathbf{X}$. Clearly, this implies that $(X_t)_{t=0}^{T}$ and $\mathbf{X}'$ have the same distribution.

Now, if $X_t = L$ for some $0 \le t \le T$, then the trajectory of $\mathbf{X}$ and $\mathbf{y}$ necessarily either intersect or cross. Consequently, we either have $X_1 \neq 0$ (which occurs with probability α) or at least one of $\mathbf{X}$ and $\mathbf{X}'$ meets $\mathbf{y}$, and is thus outside A_t for some t. By the union bound, we find

$$\begin{aligned}&\mathbb{P}\left(X_t = L \text{ for some } 0 \le t \le T \,|\, X_0 = 0\right)\\ &\qquad \le \alpha + 2\,\mathbb{P}\left(X_t \notin A_t \text{ for some } 0 \le t \le T \,|\, X_0 = 0\right).\end{aligned}$$

From (9) we know that $\mathbb{P}(X_t \le L-1 \text{ for all } 0 \le t \le T \,|\, X_0 = 0)$ is bounded away from 1 as T tends to infinity with our choice of $\alpha = \Theta(\frac{1}{T})$. So choosing $\varepsilon > 0$ small, T large enough such that

$$\mathbb{P}\left(X_t = L \text{ for some } 0 \le t \le T \,|\, X_0 = 0\right) \ge \alpha + 2\varepsilon$$

will do the job. □

Proposition 6.6
Let $\mathcal{A} = (A_t)_{t=0}^{T}$ be any separating sequence such that $0 \in A_t$ and $L \notin A_t$ for all $0 \le t \le T$. For any $0 \le a \le T$, we define the separating sequence $\mathcal{A}^a = (A^a_t)_{t=0}^{T}$ by $A^a_t := A_{t+a \,(\mathrm{mod}\ T+1)}$. Then

$$\mathcal{S}_T^{\mathcal{A}}(0, L) \le \mathcal{S}_T^{\mathcal{A}^a}(0, L) + 12\,L\alpha.$$

Proof: The case where $a = 0$ is obvious, so we may assume $a > 0$. By part (a) of Lemma 6.4, for any fixed $0 \le t \le T$, the probability that $X_t \in A_t \setminus \{0\}$ given $X_0 = 0$ is at most $2\,L\alpha$. From this we can infer

$$\begin{aligned}&\mathbb{P}\left(X_t \in A_t \text{ for all } 0 \le t \le T \,|\, X_0 = 0\right)\\ &\quad \le \mathbb{P}\left(X_t \in A_t \text{ for all } 0 \le t \le T \text{ and } X_{a-1} = X_a = X_T = 0 \,|\, X_0 = 0\right) + 6\,L\alpha\\ &\quad = \mathbb{P}\left(X_t \in A_t \text{ for all } 0 \le t \le a-1 \text{ and } X_{a-1} = 0 \,|\, X_0 = 0\right) \cdot (1-\alpha)\\ &\qquad \cdot \mathbb{P}\left(X_t \in A_t \text{ for all } a \le t \le T \text{ and } X_T = 0 \,|\, X_a = 0\right) + 6\,L\alpha\\ &\quad = \mathbb{P}\left(X_t \in A_{t-T+a-1} \text{ for all } T-a+1 \le t \le T \text{ and } X_T = 0 \,|\, X_{T-a+1} = 0\right)\\ &\qquad \cdot (1-\alpha) \cdot \mathbb{P}\left(X_t \in A_{t+a} \text{ for all } 0 \le t \le T-a \text{ and } X_{T-a} = 0 \,|\, X_0 = 0\right)\\ &\qquad + 6\,L\alpha\\ &\quad = \mathbb{P}\left(X_t \in A^a_t \text{ for all } 0 \le t \le T \text{ and } X_{T-a} = X_{T-a+1} = X_T = 0 \,|\, X_0 = 0\right)\\ &\qquad + 6\,L\alpha\\ &\quad \le \mathbb{P}\left(X_t \in A^a_t \text{ for all } 0 \le t \le T \,|\, X_0 = 0\right) + 6\,L\alpha,\end{aligned}$$

where the first and last equality follow from the Markov property and the central one from time homogeneity of the chain.

By symmetry, the same argument works for the chain started at L and the sequence of complementary sets $(S \setminus A_t)_{t=0}^T$. □

Proposition 6.7
For any separating sequence $\mathcal{A} = (A_t)_{t=0}^T$ such that $0 \in A_t$ and $L \notin A_t$ for all $0 \le t \le T$, there exists a $k \in \{0, \dots, L\}$ such that

$$\mathcal{S}_T^{\mathcal{A}}(0, L) \le \mathcal{S}_T^{\underline{\mathbf{k}}}(0, L) + 12\, L\alpha, \tag{12}$$

where $\underline{\mathbf{k}}$ denotes the constant separating sequence whose elements are all equal to the set $\{0, 1, \dots, k\}$.

Proof: By Proposition 6.6, we have

$$\mathcal{S}_T^{\mathcal{A}}(0, L) \le \frac{1}{T+1} \sum_{a=0}^{T} \mathcal{S}_T^{\mathcal{A}^a}(0, L) + 12\, L\alpha.$$

For any $0 \le k \le L$, let us define

$$f(k) := \min_{0 \le k' \le k} \frac{|\{t : k' \in A_t\}|}{T+1}.$$

Note that the function f is decreasing, $f(0) = 1$ and $f(L) = 0$. To simplify the notation, we additionally set $f(L+1) := 0$ and write $M := \max_{0 \le t \le T} X_t$. Considering only the first summands coming from each of the separation terms $\mathcal{S}_T^{\mathcal{A}^a}(0, L)$, $0 \le a \le T$, we find

$$\begin{aligned}
&\frac{1}{T+1} \sum_{a=0}^{T} \mathbb{P}\left(X_t \in A_t^a \text{ for all } 0 \le t \le T \,|\, X_0 = 0\right) \\
&\quad = \frac{1}{T+1} \sum_{a=0}^{T} \sum_{k=0}^{L} \mathbb{P}\left(M = k \,|\, X_0 = 0\right) \\
&\qquad \cdot \mathbb{P}\left(X_t \in A_t^a \text{ for all } 0 \le t \le T \,|\, M = k \text{ and } X_0 = 0\right)
\end{aligned}$$

Fix $k \in \{0, \dots, L\}$, pick k' to minimize $|\{t : k' \in A_t\}|$ over $\{0, \dots, k\}$ and note that this implies

$$f(k) = \frac{|\{t : k' \in A_t\}|}{T+1} = \frac{1}{T+1} \sum_{a=0}^{T} \mathbb{1}_{\{k' \in A_t^a\}}.$$

Let $\tau \geq 0$ be the first time when the Markov chain visits state k'. Given $M = k$, we know $\tau \leq T$, hence

$$\begin{aligned}
&\frac{1}{T+1} \sum_{a=0}^{T} \mathbb{P}\left(X_t \in A_t^a \text{ for all } 0 \leq t \leq T \,|\, M = k \text{ and } X_0 = 0\right) \\
&\quad \leq \frac{1}{T+1} \sum_{a=0}^{T} \mathbb{P}\left(k' \in A_\tau^a \,|\, M = k \text{ and } X_0 = 0\right) \\
&\quad = \frac{1}{T+1} \sum_{a=0}^{T} \sum_{t=0}^{T} \mathbb{P}\left(\tau = t \,|\, M = k \text{ and } X_0 = 0\right) \cdot \mathbb{1}_{\{k' \in A_t^a\}} \\
&\quad = \sum_{t=0}^{T} \mathbb{P}\left(\tau = t \,|\, M = k \text{ and } X_0 = 0\right) \cdot f(k) = f(k).
\end{aligned}$$

We conclude that

$$\begin{aligned}
&\frac{1}{T+1} \sum_{a=0}^{T} \mathbb{P}\left(X_t \in A_t^a \text{ for all } 0 \leq t \leq T \,|\, X_0 = 0\right) \\
&\quad \leq \sum_{k=0}^{L} \mathbb{P}\left(\max_{0 \leq t \leq T} X_t = k \,\middle|\, X_0 = 0\right) \cdot f(k) \\
&\quad = \sum_{k=0}^{L} \mathbb{P}\left(\max_{0 \leq t \leq T} X_t \leq k \,\middle|\, X_0 = 0\right) \cdot \big(f(k) - f(k+1)\big).
\end{aligned}$$

Arguing analogously in the case of $X_0 = L$, with the modifications that we consider $\min_{0 \leq t \leq T} X_t$ instead of $\max_{0 \leq t \leq T} X_t$ and τ now denotes the first time when the chain is in state k, yields

$$\begin{aligned}
&\frac{1}{T+1} \sum_{a=0}^{T} \mathbb{P}\left(X_t \notin A_t^a \text{ for all } 0 \leq t \leq T \,|\, X_0 = L\right) \\
&\quad \leq \sum_{k=0}^{L} \mathbb{P}\left(\min_{0 \leq t \leq T} X_t = k \,\middle|\, X_0 = L\right) \cdot \big(1 - f(k)\big) \\
&\quad = 1 - \sum_{k=0}^{L} \mathbb{P}\left(\min_{0 \leq t \leq T} X_t \leq k \,\middle|\, X_0 = L\right) \cdot \big(f(k) - f(k+1)\big). \\
&\quad = \sum_{k=0}^{L} \mathbb{P}\left(\min_{0 \leq t \leq T} X_t > k \,\middle|\, X_0 = L\right) \cdot \big(f(k) - f(k+1)\big),
\end{aligned}$$

where we used $1 - f(k) \geq 1 - \frac{1}{T+1} \cdot |\{t : k \in A_t\}|$ to derive the inequality and the last equality follows from $\sum_{k=0}^{L} f(k) - f(k+1) = 1$. By combining these two estimates, it follows that

$$\frac{1}{T+1} \sum_{a=0}^{T} \mathcal{S}_T^{A^a}(0, L) \leq \sum_{k=0}^{L} \big(f(k) - f(k+1)\big) \cdot \mathcal{S}_T^{\mathbf{k}}(0, L).$$

From the fact that the coefficients $f(k) - f(k+1),\ 0 \le k \le L$, sum up to 1, plus $f(L) = f(L+1)$, we can conclude that there exists some $k \in \{0, \dots, L-1\}$ such that

$$\frac{1}{T+1} \sum_{a=0}^{T} \mathcal{S}_T^{\mathcal{A}^a}(0, L) \le \mathcal{S}_T^{\mathbf{k}}(0, L),$$

which completes the proof. □

Combining Propositions 6.5 and 6.7, it follows that for any fixed $L \ge 1$, any $\alpha = \alpha(T) = \Theta\left(\frac{1}{T}\right)$ and T sufficiently large, the optimal separation $\mathcal{S}_T(0, L)$ is the maximum of 1 and

$$\mathcal{S}_T^{\mathbf{k}}(0, L) + O(L\alpha) \le e^{-\alpha T/(k+1)} + e^{-\alpha T/(L-k)} + O(L\alpha), \tag{13}$$

for $0 \le k < L$, where the inequality follows from (9).

To finish the proof of Theorem 6.2, we need one more elementary estimation:

Lemma 6.8
Given $A > 0$, define the function

$$f_A(x) := e^{-A/x} + e^{-A/(1-x)}, \quad x \in (0, 1).$$

Then it holds

$$\sup_{0<x<1} f_A(x) = \max\left(e^{-A}, 2\, e^{-2A}\right).$$

PROOF: First note that the function f_A lies in $C^\infty((0,1))$ and is symmetric around $x = \frac{1}{2}$. Calculating its first two derivatives gives

$$f_A'(x) = \tfrac{A}{x^2} \cdot \mathrm{e}^{-\frac{A}{x}} - \tfrac{A}{(1-x)^2} \cdot \mathrm{e}^{-\frac{A}{1-x}} \quad \text{and}$$
$$f_A''(x) = \left(\tfrac{A^2}{x^4} - \tfrac{2A}{x^3}\right) \cdot \mathrm{e}^{-\frac{A}{x}} + \left(\tfrac{A^2}{(1-x)^4} - \tfrac{2A}{(1-x)^3}\right) \cdot \mathrm{e}^{-\frac{A}{(1-x)}}.$$

If $x \in (0,1)$ is a stationary point of f_A, we necessarily have

$$\tfrac{A}{x^2} \cdot \mathrm{e}^{-\frac{A}{x}} = \tfrac{A}{(1-x)^2} \cdot \mathrm{e}^{-\frac{A}{1-x}}$$

and as a consequence the sign of $f_A''(x)$ is given by the sign of

$$g_A(x) := 2\,(A+1)\,x^2 - 2\,(A+1)\,x + A.$$

Due to $A > 0$, the function g_A is strictly convex. Assuming the existence of two local maxima of f_A on $(0,1)$ – at points $x_1 < x_2$ say – forces the existence of a local minimum at $x_3 \in (x_1, x_2)$. Hence $\max\{f_A''(x_1), f_A''(x_2)\} \le 0 \le f_A''(x_3)$, which contradicts the strict convexity of g_A.

Consequently, f_A can have at most one local maximum in $(0,1)$, which then lies at $x = \frac{1}{2}$ for symmetry reasons. In conclusion, f_A either attains its maximum on $(0,1)$ at $x = \frac{1}{2}$ or converges to its supremum on the boundary. □

PROOF OF THEOREM 6.2: Applying Lemma 6.8 to (13) with $A = \alpha T/(L+1)$, we see that we can pick $\delta > 0$ arbitrarily small, choose

$$\alpha = \alpha(T) = \frac{1}{2}\left(\ln(2) + \delta\right) \cdot \frac{L+1}{T}$$

and find that the optimal separation of the states 0 and L in the chain $\mathbf{X}$ on $[0, T]$ is 1 if T is chosen sufficiently large (after having fixed L). Hence, by Theorem 6.3 and (8) we know that the chain $(X_t)_{t=0}^{T}$ segregates 0 and L. From (8) we can further read off that for L fixed

$$\|P^T(0,.) - P^T(L,.)\|_{\mathrm{TV}} = \mathrm{e}^{-\left(\ln(2)+\delta\right)\frac{L+1}{L}} - o_T(1). \tag{14}$$

Given $\varepsilon > 0$, we can choose $\delta > 0$ small, L large enough and then pick T sufficiently large to make the right-hand side of (14) larger than $\frac{1}{2} - \varepsilon$. This completes the proof. □

Remark 2
One downside of the implicit construction proving Theorem 6.2 is the fact that it does not give much information about the coupling involved. As the coupling will have to take into account the whole trajectories of the two individual copies, it is highly unlikely that the coupled process will have the Markov property. In this respect, it is still an open problem if the value of $\frac{1}{\mathrm{e}}$, established in Section 5, can be pushed further (as supremum of achievable total variation distances that can be retained in segregating Markov chains), if we restrict ourselves to *Markovian* couplings.

Acknowledgements

First of all, we want to thank Jeff Steif for bringing this peculiar kind of Markov chains to our attention. Then, we want to give thanks to him, Olle Häggström and Peter Hegarty for their valuable comments on earlier drafts.

References

[1] BOLLOBÁS, B., "Modern Graph Theory", Springer, 1998.

[2] FREEMAN, D., "Markov Chains", Springer, 1983.

[3] HÄGGSTRÖM, O., A Note on Disagreement Percolation, *Random Structures & Algorithms*, Vol. 18 (3), pp. 267-278, 2001.

[4] LEVIN, D.A., PERES, Y. and WILMER, E.L., *"Markov Chains and Mixing Times"*, American Mathematical Society, 2008.

[5] Rosenthal, J.S., Faithful Couplings of Markov Chains: Now Equals Forever, *Advances in Applied Mathematics*, Vol. 18 (3), pp. 372-381, 1997.

Timo Hirscher
Department of Mathematical Sciences,
Chalmers University of Technology
and University of Gothenburg,
412 96 Gothenburg, Sweden.
hirscher@chalmers.se

Anders Martinsson
Department of Mathematical Sciences,
Chalmers University of Technology
and University of Gothenburg,
412 96 Gothenburg, Sweden.
andemar@chalmers.se